20.7 Circuit Techniques to Enable 430Gb/s/mm² Proximity Communication

David Hopkins, Alex Chow, Robert Bosnyak, Bill Coates, Jo Ebergen, Scott Fairbanks, Jon Gainsley, Ron Ho, Jon Lexau, Frankie Liu, Tarik Ono, Justin Schauer, Ivan Sutherland, Robert Drost

Sun Microsystems Laboratories, Menlo Park, CA

Proximity Communication (PxIO) is an interchip communication technology that offers important advantages over conventional chip-to-chip I/O, including high channel density and low power. Proximity Communication places chips face-to-face, within a few microns of one another [1]. Capacitive coupling between top-level metal I/O structures permits efficient communication.

Previous reports of chip-to-chip I/O based on electromagnetic coupling have demonstrated performance improvements in channel density, power, or throughput over conventional I/O [2-4]. However, they do not address issues with chip alignment in order to maintain reliable communication in a production environment, where transmitting and receiving channels may become misaligned due to vibrations or thermal expansion of chips. Also, inductively-coupled implementations are prone to crosstalk noise [2]. Novel techniques are presented here to mitigate misalignment and crosstalk, and shown to enable reliable and power efficient Proximity Communication.

The main challenge with high pad density is the mechanical alignment accuracy it requires. To relax this constraint, an *electronic alignment correction* technique [5] has been developed. This technique shifts the location of the transmitting channel to compensate for physical misalignment between the transmitting and receiving channels. Each transmit pad is physically divided into a 4×4 array of micropads (Fig. 20.7.1). Multiplexers on the transmitting chip steer data to the micropads that best align with the receiving chip. The optimal multiplexer configuration is determined by precisely measuring chip alignment.

Electronic alignment increases power consumption due to the extra wires and multiplexers necessary for data steering. To reduce this power cost, a power-efficient multiplexer is designed that uses NMOS-only pass gates (Fig. 20.7.1). A low select signal drives M2's gate low, making it opaque. A high select signal drives M2's gate to one threshold voltage below V_{DD}. It is held very weakly in this state, as M1 is off. A rising data transition bootstraps M2's gate above V_{DD}, allowing M2 to pass full V_{DD} levels. Falling transitions restore the gate voltage on M2 to one threshold voltage below V_{DD}. Because M2's gate voltage tracks the channel voltage, the effective channel capacitance and resistance is lower. Compared to a typical CMOS pass gate, the bootstrapped NMOS pass gate reduces overall transmitter power by more than 20%, from 2.5pJ/b to 2.0pJ/b, while providing the same edge rates. Also, since the multiplexers consist of only NMOS devices, layout is more compact. This technique is used in memory design and demonstrates both high performance and reliability [6]. However, it requires occasional data transitions to prevent droop and jitter if the data remains high over an extended time on the order of 1ms.

Another performance limiter for high density I/O technologies is crosstalk noise [2]. An arrangement of channels is developed called *butterfly differential signaling* (Fig. 20.7.2) [7], which completely rejects nearest neighbor crosstalk, for a receiver with good common mode rejection. Figure 20.7.2 highlights a differential channel (pads A+ and A-) and four adjacent channels (B,C,D,E). Channel A sees no net noise from channel B, because pads B+ and B- couple equally to A+; any noise due to a transition on B+ is canceled by an opposing transition on B-. E+ and E- act similarly on A-. Channel A sees no net noise from D or C, because D+ and C- couple equally to A+ and A-; any noise due to a transition on D or C is thus common-mode to A. This crosstalk cancellation scheme enables reliable communication using smaller I/O pads, over a greater chip separation, and at higher data rates. This pad arrangement can also mitigate crosstalk in any 2D array of communication channels, including channels on different layers on a PCB.

For ease of scaling, the PxIO circuitry has been designed as modular slices that can be easily tiled and assembled (Fig. 20.7.2). Each slice consists of 36 data and 8 timing channels. Each channel consists of a pair of differential receiver pads on a pitch of 36μm. Although two timing channels per slice are sufficient for bidirectional I/O, 8 channels have been allocated to accommodate more complex timing protocols such as asynchronous handshaking. The butterfly signaling layout pattern and the need for electronic alignment across slice boundaries complicate the design of these slices. The boundary of each slice is jagged, allowing a slice to contain only complete pairs of differential pads.

Because signals are capacitively-coupled, explicit DC-biasing is necessary on the receiving plates. Disabled PMOS transistors are used to form a large resistance between each amplifier input and the applied bias voltages (Fig. 20.7.3). The receiver for a data channel is a simple clocked sense amplifier, while that of a timing channel is a resistively-loaded source coupled pair.

A test chip was fabricated in a TSMC 0.18μm CMOS process. The chip has four PxIO slices for a total of 72 transmit and 72 receive data channels. All channels operate simultaneously at 1.8Gb/s per channel. This provides an aggregate I/O bandwidth of 260Gb/s on each chip, equivalent to 430Gb/s/mm². On PRBS data, the measured bit error rate (BER) is lower than 10^{-15}. The combined energy cost of the transmitter, receiver and amortized clock distribution is 3.0pJ/b; this is comparable to the lowest previously achieved [2]. However, the present implementation offers electronic alignment correction that provides much greater tolerance to misalignment and the noise cancellation scheme results in a BER that is 1000× lower.

Figure 20.7.4 shows the voltage and timing waterfall curves. At 1.6Gb/s and a BER below 10^{-14}, the link shows a 0.72 UI timing margin and more than 200mV of voltage margin. Figure 20.7.5 shows the received data eye for a BER of 10^{-11} and 1.8Gb/s per channel, determined by measuring the timing margin at different voltage offsets. These offsets are intentionally introduced by applying a different DC-bias voltage to each of the inputs of the amplifier. The wide eye opening indicates that the maximum per-channel data rate is limited by supporting circuitry, and not the PxIO channel itself. Figure 20.7.6 demonstrates that the performance of PxIO degrades as the interchip distance is increased. No bit error is measured until the interchip separation exceeds 9μm. Applying receiver offset compensation can extend this range. All measurements are taken with air as the interchip dielectric; this tolerance can be significantly increased by using interposer materials with higher permittivity.

Acknowledgments:
We thank P. Bell-Bosnyak, A. Chakraborty, K. Dhaliwal, P. Teehan, and M. Yoo for help in chip design and layout. This work was supported in part by DARPA contract NBCH30390002.

References:
[1] R. Drost, et al., "Proximity Communication," *IEEE J. Solid-State Circuits*, vol. 39, pp. 1529-1535, Sept., 2004.
[2] N. Miura, et al., "A 1Tb/s 3W Inductive-Coupling Transceiver for Inter-Chip Clock and Data Link," *ISSCC Dig. Tech. Papers*, pp. 424-425, Feb., 2006.
[3] K. Kanda, et al., "1.27 Gb/s/pin 3mW/pin Wireless Superconnect (WSC) Interface Scheme," *ISSCC Dig. Tech. Papers*, pp. 186-187, Feb., 2003.
[4] S. Mick, J. Wilson, and P. Franzon, "4Gbps High-Density AC Coupled Interconnection," *IEEE Proc. CICC*, pp. 133-140, 2002.
[5] R. Drost, et al., "Electronic Alignment for Proximity Communication," *ISSCC Dig. Tech. Papers*, pp. 144-145, Feb., 2004.
[6] R. J. Proebsting, technical discussions.
[7] R. Drost, "Butterfly Differential Signaling," US Patent pending, filed Sept. 12, 2003.

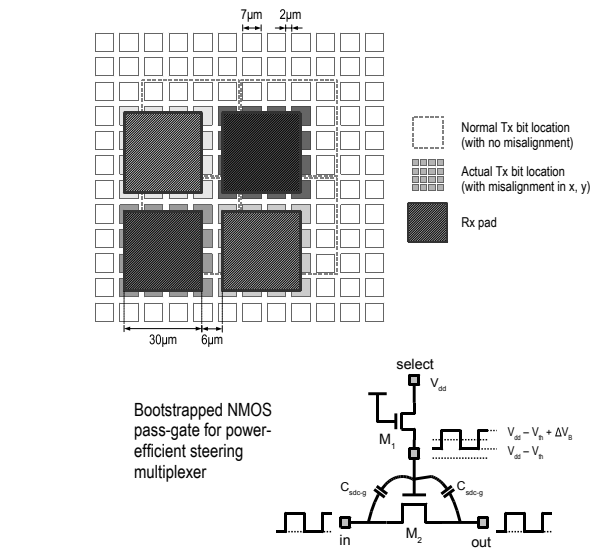

Figure 20.7.1: Data steering using boostrapped NMOS.

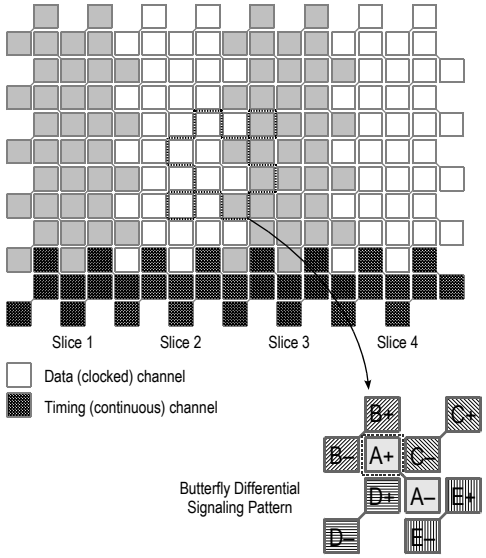

Figure 20.7.2: Modular butterfly PxIO array.

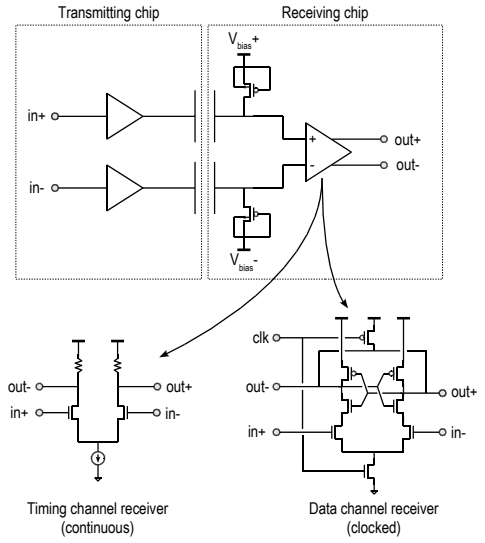

Figure 20.7.3: Receiver circuits for timing and data.

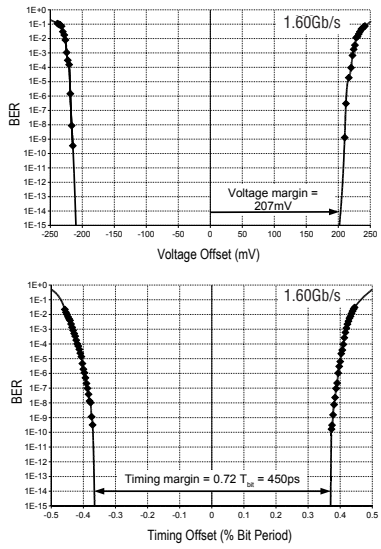

Figure 20.7.4: Timing/voltage waterfall curves.

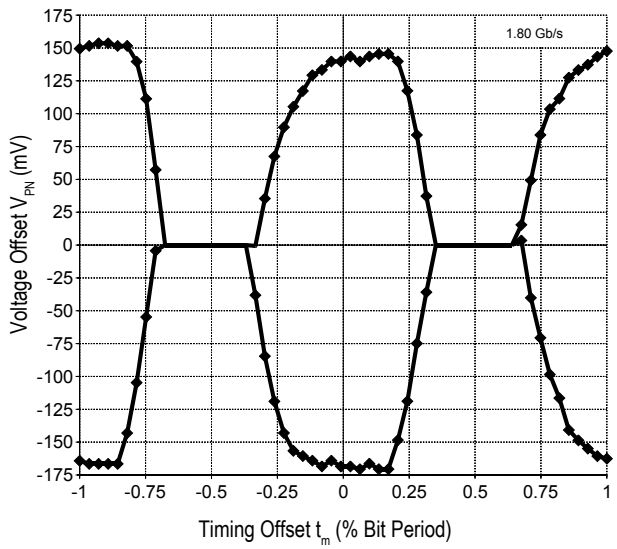

Figure 20.7.5: Received signal eye opening.

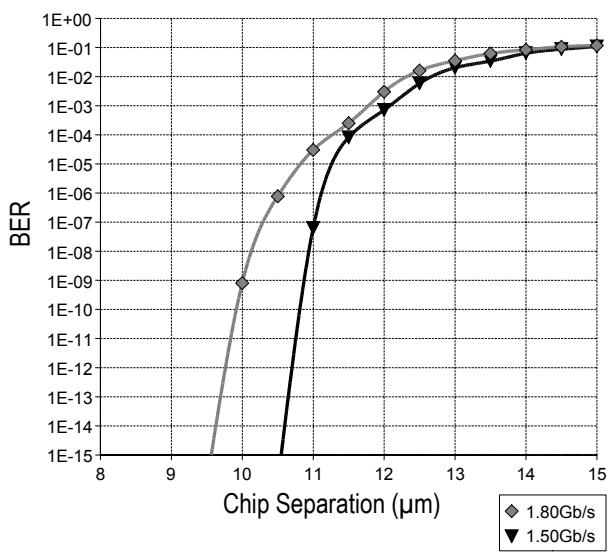

Figure 20.7.6: Measured BER vs. chip separation.

Continued on Page 609

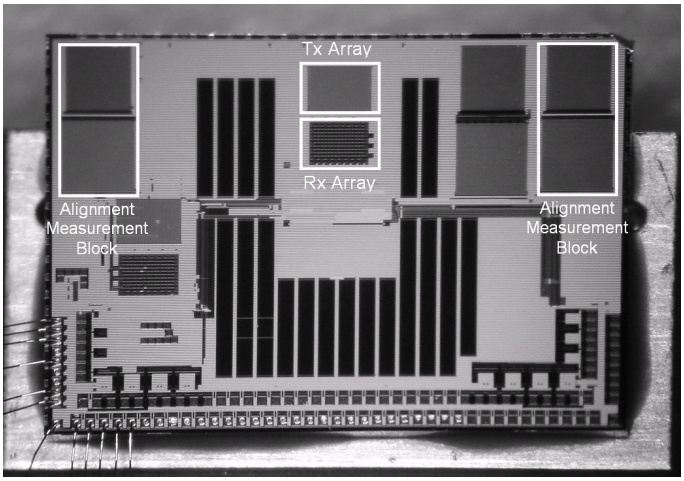

Figure 20.7.7: Die micrograph.

